From Bottleneck to Benchmark: Using AI to Automate the Requirements Traceability Matrix (RTM) for FDA Submissions

Author: Shay Ginsbourg, M.Sc.

Date: January 20, 2026

1. Executive Summary

In the highly regulated landscape of medical device development, the Requirements Traceability Matrix (RTM) stands as a cornerstone of regulatory compliance and product safety. For decades, the RTM has been the primary mechanism for demonstrating that every design input is accounted for in the final product and that every safety requirement has been rigorously verified and validated. However, as medical devices increasingly rely on complex software and iterative development cycles, the traditional manual approach to maintaining an RTM has become a significant bottleneck. This whitepaper explores the transformative potential of Artificial Intelligence (AI) in automating the RTM process, shifting it from a labor-intensive compliance burden to a high-performance benchmark for engineering excellence.

The “bottleneck” of manual traceability is characterized by fragmented data, human error, and a reactive approach to documentation that often lags behind actual development. By leveraging advanced AI techniques such as Natural Language Processing (NLP), Large Language Models (LLMs), and semantic mapping, manufacturers can now automate the extraction, linking, and maintenance of traceability artifacts. This shift not only significantly reduces documentation time but also enhances the accuracy and completeness of FDA submissions. Through real-world case studies, including the success of Meta Reality Labs, we demonstrate how AI-driven traceability enables faster release cycles while maintaining the highest standards of safety and regulatory adherence.

2. Introduction: The Regulatory Landscape

The necessity of traceability in medical device development is not merely a best practice; it is a legal requirement mandated by the U.S. Food and Drug Administration (FDA) under 21 CFR 820.30, also known as Design Controls. These regulations require manufacturers to establish and maintain procedures to ensure that the design requirements are appropriate and address the intended use of the device, including the needs of the user and patient. Specifically, the FDA emphasizes the linkage between design inputs, design outputs, design verification, and design validation.

Regulatory Element	Description	RTM Role
Design Input	Physical and performance requirements of a device.	The starting point of the trace.
Design Output	The results of a design effort at each design phase.	Linked to inputs to show fulfillment.
Design Verification	Confirmation that design output meets design input.	Linked to test cases and results.
Design Validation	Confirmation that device specifications conform to user needs.	Linked to clinical or user evaluations.

Beyond the FDA, international standards such as ISO 13485:2016 and IEC 62304 further solidify the requirement for traceability. ISO 13485 requires a quality management system that ensures product safety and efficacy through documented evidence of design and development. IEC 62304, which focuses on medical device software lifecycle processes, explicitly requires traceability between software requirements, software system tests, and risk control measures. As software-driven medical devices (SaMD) become more prevalent, the volume of requirements and the frequency of changes have grown exponentially, making traditional traceability methods increasingly untenable.

3. The Traceability Bottleneck: Current Industry Challenges

Despite its critical importance, the creation and maintenance of an RTM is often cited as one of the most arduous tasks in the development lifecycle. The "Traceability Bottleneck" arises from several systemic challenges that plague traditional medical device manufacturing:

Manual Processes and Human Error

Most RTMs are still managed using spreadsheets or static documents. This manual entry is inherently prone to human error, such as missed links, incorrect references, or outdated information. In a system with thousands of requirements, even a small percentage of errors can lead to significant regulatory risks, including FDA warning letters or product recalls.

Tool Fragmentation

Modern development teams use a variety of specialized tools: Jira for requirements management, GitHub or GitLab for version control, and TestRail or Zephyr for test management. These tools often operate in silos, making it difficult to maintain a unified view of traceability. Manually bridging these gaps requires constant cross-referencing and data duplication, which consumes valuable engineering resources.

The "Living Document" Problem

In an Agile or iterative development environment, requirements are constantly evolving. A single change in a high-level user need can ripple through system requirements, design specifications, and dozens of test cases. Keeping a manual RTM up-to-date in real-time is nearly impossible, leading to a "documentation lag" where the RTM only reflects the state of the product weeks or months after the code has been written.

Impact on Timelines and Costs

The overhead of manual traceability is staggering. Industry research by McKinsey & Company suggests that quality assurance and regulatory compliance can consume 12-18% of total industry revenue. Furthermore, Gartner reports that project costs increase by 30-40% when formal validation is required. These costs are not just financial; they represent a significant delay in bringing life-saving technologies to patients.

4. The AI Revolution in Requirements Management

The emergence of Artificial Intelligence, particularly in the realm of Generative AI and Natural Language Processing (NLP), offers a paradigm shift in how requirements are managed and traced. Rather than viewing the RTM as a static artifact to be completed at the end of a project, AI enables a "traceability-by-design" approach where links are generated and maintained as a natural byproduct of the development process.

Natural Language Processing (NLP) for Requirement Extraction

One of the first hurdles in RTM automation is the extraction of requirements from unstructured or semi-structured documents. NLP algorithms can be trained to identify "shall" statements, safety-critical constraints, and performance metrics within vast technical specifications. This automated extraction ensures that no requirement is overlooked and provides a structured foundation for the traceability matrix.

Large Language Models (LLMs) for Semantic Mapping

Traditional traceability tools rely on exact keyword matching or manual ID linking. However, requirements and test cases are often written by different teams using different terminology. LLMs excel at understanding the semantic intent behind the text. For example, an LLM can recognize that a requirement for "battery longevity" is directly related to a test case for "power consumption under peak load," even if the specific words do not overlap. This semantic mapping significantly reduces the manual effort required to establish initial links.

Retrieval-Augmented Generation (RAG) for Contextual Accuracy

To ensure that AI-generated links are accurate and relevant to a specific project, Retrieval-Augmented Generation (RAG) is employed. RAG allows an AI system to reference a project's own internal documentation, Quality Management System (QMS) procedures, and historical data when suggesting traceability links. This ensures that the AI's recommendations are grounded in the specific context of the device being developed, rather than general industry knowledge.

AI Agents: From Passive Tools to Proactive Teammates

The most advanced application of AI in RTM automation is the use of AI Agents. Unlike traditional software that waits for user input, AI agents proactively monitor the development toolchain. When a developer updates a requirement in Jira or pushes a new test script to GitHub, the AI agent automatically analyzes the change, identifies the impacted links in the RTM, and suggests updates. This shifts the role of the quality engineer from a "data entry clerk" to a "reviewer," significantly accelerating the compliance workflow.

5. Automating the RTM: A Technical Deep Dive

Automating the RTM requires a sophisticated integration of several AI technologies. The following table outlines the core technical components of an AI-driven traceability system:

AI Technology	Application in RTM	Benefit
Vector Embeddings	Converting text into mathematical vectors to measure similarity.	Enables high-accuracy semantic mapping between artifacts.
Change Impact Analysis	Automatically identifying all downstream artifacts affected by a change.	Ensures the RTM remains consistent and up-to-date in real-time.
Redundancy Detection	Identifying duplicate or overlapping requirements.	Streamlines the RTM and reduces testing overhead.
Automated Test Generation	Using LLMs to draft test cases based on requirement specifications.	Ensures 100% test coverage and reduces manual drafting time.

Semantic Vector Embeddings

At the heart of AI-driven traceability is the use of vector embeddings. By representing requirements and test cases as high-dimensional vectors, the system can calculate the

“cosine similarity” between them. This allows the system to suggest links with a high degree of confidence, even across different levels of abstraction (e.g., linking a high-level user need to a low-level software unit test).

Automated Change Impact Analysis

One of the most powerful features of AI automation is the ability to perform real-time change impact analysis. When a requirement is modified, the AI agent scans the entire traceability graph to identify every linked design specification, risk control, and test case. It then flags these items for review, ensuring that no change is made in isolation and that the entire system remains validated.

Redundancy and Conflict Detection

In large-scale projects, it is common for requirements to become redundant or even contradictory over time. AI agents can analyze the entire requirement set to identify overlaps or conflicts. For instance, if one requirement specifies a maximum weight of 5kg and another specifies a minimum weight of 6kg, the AI will flag this conflict immediately, allowing the team to resolve it before it impacts the design or testing phases.

5.1 Deep Dive: The Role of Semantic Vector Embeddings in Traceability Recovery

The technical foundation of AI-driven traceability recovery lies in the transformation of natural language artifacts into high-dimensional mathematical representations known as vector embeddings. In a traditional RTM, a human auditor might look for keyword overlaps—for example, ensuring that a requirement mentioning “heart rate” is linked to a test case also containing the phrase “heart rate.” However, this approach fails when different teams use synonymous but distinct terminology, such as “cardiac frequency” or “pulse monitoring.”

Semantic vector embeddings solve this problem by mapping words and sentences into a continuous vector space where the distance between vectors represents the semantic similarity between the underlying concepts. Modern embedding models, such as those based on the Transformer architecture, are trained on billions of lines of text, allowing them to understand that “heart rate” and “cardiac frequency” are conceptually identical.

Feature	Traditional Keyword Matching	AI Semantic Mapping
Matching Logic	Exact string or substring match.	Conceptual and contextual similarity.
Handling Synonyms	Requires manual mapping or extensive thesauri.	Inherently understand synonyms and related terms.
Context Awareness	Limited; often misses context-dependent links.	High; understands the role of a term within a sentence.
Scalability	Decreases as the number of artifacts grows.	Increases: AI can process thousands of links in seconds.

When applied to an FDA submission, this technology allows for the automatic recovery of missing links in a legacy Design History File (DHF). By scanning thousands of requirements and test cases, the AI can identify potential gaps where a requirement exists but no corresponding test case has been linked. This “traceability recovery” is a critical step for manufacturers preparing for an audit or a new 510(k) submission, as it ensures that the evidence provided to the FDA is both complete and robust.

5.2 Automated Change Impact Analysis: Ensuring Consistency in Agile Environments

One of the most significant advantages of AI-driven RTMs is their ability to handle the rapid pace of change in modern software development. In an Agile environment, requirements are not static; they evolve as the product is refined through successive sprints. In a manual system, a single change to a high-level requirement can necessitate hours of manual work to identify and update all downstream artifacts, including system requirements, design specifications, risk control measures, and verification tests.

AI agents automate this process through continuous monitoring of the development toolchain. When a change is detected in a source artifact (e.g., a requirement in Jira), the AI agent performs a “graph traversal” of the traceability matrix to identify all directly and indirectly linked items.

1. **Detection**: The AI agent identifies a modification to a requirement.

2. **Analysis**: The agent calculates the "blast radius" of the change, identifying all affected downstream artifacts.

3. **Recommendation**: The agent suggests specific updates to the linked artifacts to maintain consistency.

4. **Verification**: The agent flags the affected test cases as "out-of-date," requiring them to be re-run to ensure the change has not introduced new defects.

This automated workflow ensures that the RTM is always a "living document" that accurately reflects the current state of the product. For FDA submissions, this means that the manufacturer can provide a real-time snapshot of their traceability at any point in the development lifecycle, rather than a static document that was manually compiled at the end of the project. This level of transparency and consistency is highly valued by regulatory reviewers, as it demonstrates a high degree of control over the design process.

8.1 Overcoming the "Black Box" Challenge: Explainable AI (XAI) in Compliance

A significant hurdle in the adoption of AI for regulatory documentation is the perceived "black box" nature of many AI models. Regulatory bodies like the FDA require a high degree of transparency and explainability, especially when AI is used to make decisions that impact product safety. To address this, modern AI-driven compliance tools are incorporating Explainable AI (XAI) techniques. XAI ensures that when an AI agent suggests a traceability link or identifies a conflict, it provides a clear rationale for its decision. For example, instead of simply stating that "Requirement A is linked to Test Case B," the AI might provide a summary of the semantic similarities it identified or highlight the specific technical parameters that connect the two artifacts.

- **Transparency**: Providing a clear audit trail of how AI-generated links were created and who reviewed them.

- **Interpretability**: Using models that allow humans to understand the underlying logic of the AI's recommendations.

- **Accountability**: Ensuring that every AI-driven action is subject to human oversight and approval.

By prioritizing explainability, manufacturers can build trust with regulatory reviewers and demonstrate that their use of AI is both responsible and controlled. This approach aligns with the FDA's guiding principles for Good Machine Learning Practice, which emphasize the need for transparency and the importance of maintaining human accountability in the development of AI-enabled medical products.

6. Case Study: Benchmarking Success

The theoretical benefits of AI-driven RTM automation are best illustrated through real-world applications. A prominent example is Meta Reality Labs, Meta's hardware and AI division, which has been at the forefront of developing regulated medical products, such as biosensing wearables and Software as a Medical Device (SaMD), including AI-enabled smart glasses.

The Challenge

Meta Reality Labs faced the daunting task of maintaining rigorous regulatory compliance while operating at the high velocity typical of a leading technology company. Key challenges included preserving engineering velocity amid design controls, eliminating manual documentation overhead, and achieving real-time traceability to ensure products were always audit-ready without disrupting rapid iteration or developer workflows. Traditional manual traceability methods were incompatible with their rapid development cycles and would have significantly delayed their time-to-market.

The Solution

By implementing an AI-driven compliance platform (Ketryx), Meta was able to automate the creation and maintenance of their RTM. The system integrated directly with their existing engineering tools (Jira, GitHub, and internal tools), providing automated evidence generation, embedded compliance controls, real-time traceability, safety-critical prioritization, and AI-assisted reviews. This unified, AI-native approach embedded quality and regulatory controls directly into engineering workflows, eliminating the need for parallel systems.

The Results

The impact of this automation was transformative, setting a new benchmark for the industry:

- **Reduced Review Cycles from Weeks to Minutes**: AI-assisted reviews streamlined compliance checks, dramatically accelerating the process.

- **Sustained Engineering Velocity**: The platform enabled continued use of Meta's CI/CD-grade development practices, allowing teams to make frequent, safe changes without slowing innovation.

- **Automated Documentation and Continuous Readiness**: Continuous traceability and automated checks supported ongoing regulatory readiness, increasing confidence in product safety and compliance.

- **Scalable Compliance Foundation**: The system established a foundation that scales with products, teams, and ambitions, enabling faster, safer innovation at scale.

As Lucas Fernandez, Director of Medical Devices at Meta, stated: "Ketryx AI is the most scalable approach we've used for regulated product development. It lets our teams make frequent, safe changes while staying continuously audit-ready, without relying on linear headcount growth or stopping teams for compliance exercises."

Improved Engineering Culture

Developers were able to stay within their preferred tools while the AI handled the compliance overhead in the background.

7. The FDA's Perspective on AI in Documentation

A common concern among medical device manufacturers is whether the FDA will accept AI-generated documentation. The FDA's stance has evolved significantly with the establishment of the Digital Health Center of Excellence and the publication of several guiding documents regarding AI and Machine Learning (AI/ML).

Fostering Responsible Innovation

The FDA has explicitly stated its goal to foster responsible and high-quality digital health innovation. In its "AI/ML Software as a Medical Device Action Plan," the agency outlines a tailored regulatory framework that supports the iterative nature of AI-driven development. While the FDA is primarily focused on AI within the device itself, the same principles of transparency, validation, and risk management apply to the use of AI in the development process.

The Importance of "Human-in-the-Loop"

A critical requirement for the acceptance of AI-driven RTMs is the "Human-in-the-Loop" (HITL) principle. The FDA expects that AI will be used to assist human experts, not replace them. In an automated RTM workflow, the AI suggests links and identifies gaps, but a qualified human professional must review and approve these actions. This ensures that the final RTM remains a product of human judgment and accountability.

Validation of AI Tools

Manufacturers using AI to automate their RTM must also validate the AI tools themselves. This involves demonstrating that the AI system consistently produces accurate results and that its performance is monitored over time. The FDA's Good Machine Learning Practice (GMLP) guiding principles provide a roadmap for this validation, emphasizing the need for high-quality data, robust model testing, and clear documentation of the AI's decision-making process.

8. Implementation Strategy: Moving from Bottleneck to Benchmark

Transitioning from a manual RTM to an AI-automated system requires a strategic approach. Organizations should follow a structured roadmap to ensure successful implementation:

Phase 1: Assessment and Maturity Mapping

Evaluate the current state of traceability within the organization. Identify the primary pain points, such as tool silos, high error rates, or slow release cycles. Determine the organization's "traceability maturity" and set clear goals for automation.

Phase 2: Tool Selection and Integration

Select an AI-driven traceability platform that integrates seamlessly with the existing engineering stack (e.g., Jira, GitHub, Azure DevOps). Ensure the tool supports semantic mapping, automated change impact analysis, and HITL workflows.

Phase 3: Pilot Program

Implement the AI system on a single, well-defined project. This allows the team to refine the AI's configuration, establish validation protocols, and demonstrate early wins to stakeholders.

Phase 4: Scaling and Cultural Shift

Once the pilot is successful, scale the system across the organization. This phase requires a cultural shift, as engineers and quality professionals must learn to trust and collaborate with AI agents. Continuous training and clear communication are essential to overcome resistance to change.

Implementation Step	Key Activity	Expected Outcome
Data Cleaning	Standardizing requirement formats and terminology.	Improved AI accuracy and mapping.
Workflow Design	Defining the human review and approval process.	Regulatory compliance and accountability.
Validation	Formally validating the AI tool for its intended use.	Audit readiness and FDA acceptance.

9. Conclusion: The Future of Regulatory Compliance

The transition from a manual, bottlenecked Requirements Traceability Matrix to an AI-automated benchmark represents a fundamental shift in medical device development. As the complexity of life-saving technologies continues to grow, the industry can no longer afford the inefficiencies and risks associated with traditional documentation methods. AI-driven traceability offers a path forward that harmonizes the seemingly contradictory goals of high-velocity innovation and rigorous regulatory compliance.

By automating the arduous tasks of requirement extraction, semantic mapping, and change impact analysis, manufacturers can redirect their most valuable resources—their engineers—toward solving the complex medical challenges of the 21st century. The success of early adopters like Meta Reality Labs demonstrates that this is not a distant future, but a present-day reality. Those who embrace AI-driven compliance will not only accelerate their time-to-market but also set a new standard for product safety, reliability, and engineering excellence.

In the final analysis, the RTM should not be a hurdle to be cleared at the end of a project; it should be a strategic asset that provides real-time visibility into the integrity of the design. With AI as a proactive teammate, the medical device industry can finally move beyond the traceability bottleneck and establish a new benchmark for what is possible in regulated software development.

10. References

1. U.S. Food and Drug Administration. (2024). 21 CFR Part 820 - Quality System Regulation. https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/CFRSearch.cfm?CFRPart=820

2. International Organization for Standardization. (2016). ISO 13485:2016 - Medical devices — Quality management systems — Requirements for regulatory purposes. https://www.iso.org/standard/59752.html

3. International Electrotechnical Commission. (2006). IEC 62304:2006 - Medical device software — Software life cycle processes. https://webstore.iec.ch/publication/6848

4. Ketryx. (2023). The Ultimate Guide to Requirements Traceability Matrix (RTM). https://www.ketryx.com/blog/the-ultimate-guide-to-requirements-traceability-matrix-rtm

5. ArXiv. (2024). Natural Language Processing for Requirements Traceability. https://arxiv.org/html/2405.10845v1

6. McKinsey & Company. (2022). The future of quality in the medical device industry. https://www.mckinsey.com/industries/life-sciences/our-insights/the-future-of-quality-in-the-medical-device-industry

7. Gartner. (2022). Market Guide for Medical Device Quality Management Systems. https://www.gartner.com/en/documents/4015699

8. U.S. Food and Drug Administration. (2025). Artificial Intelligence and Machine Learning Software as a Medical Device Action Plan. https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device-action-plan

9. U.S. Food and Drug Administration. (2025). Good Machine Learning Practice for Medical Device Development: Guiding Principles. https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles

10. Ketryx. (2025). AI Agents for Regulated Software Teams. https://www.ketryx.com/capabilities/ai-agents

11. ScienceDirect. (2013). Medical device standards' requirements for traceability during the software development lifecycle. https://www.sciencedirect.com/science/article/pii/S0920548913000780

12. Nature Digital Medicine. (2024). A scoping review of reporting gaps in FDA-approved AI medical devices. https://www.nature.com/articles/s41746-024-01270-x

13. FDA Digital Health Center of Excellence. (2025). Digital Health Center of Excellence. https://www.fda.gov/medical-devices/digital-health-center-excellence

14. Ketryx Case Study. (2025). How Meta Reality Labs Accelerates Product Innovation with Ketryx. https://www.ketryx.com/case-studies/how-meta-reality-labs-accelerates-product-innovation-with-ketryx

For more information, follow me on LinkedIn at https://www.linkedin.com/in/shayginsbourg/